

Docket No.: 60751.300101

Patent

REMARKS

The Examiner is thanked for the comments in the Action. They have helped us considerably in understanding the Action and in drafting this Response thereto.

It is our understanding that claims 1-23, 25-31, and 33-40 remain pending in this application, wherein claims 24 and 32 are herein canceled and claims 23, 25, 31, and 33 are amended for reasons specifically remarked upon below.

The § 102(b) rejections, generally:

Claims 1-40 (all) are rejected as being anticipated by Yamada. Respectfully this is error.

The Examiner has correctly noted that both Yamada and the present application describe interrupt controllers for multiple processors and that both allow interrupt requests to be routed to proper processors. However, there are major differences in how the respective inventions accomplish this.

The § 102(b) rejections as applied to claims 1, 12, 23, 31 and 39-40:

At col. 3, ln. 7-31, Yamada teaches:

Only one of the interrupt right control circuits 3-1 to 3-n receives an interrupt request from an I/O interface at a timing to supply the interrupt signal (request) to a corresponding CPU. A right of receiving an interrupt request from an I/O interface to supply the interrupt signal to a CPU is called an "interrupt right". Each of the interrupt right control circuits 3-1 to 3-n having the interrupt right supplies the interrupt signal IA to an interrupt terminal INT of a corresponding one of the CPUs 2-1 to 2-n when it receives an interrupt request from the interrupt controller 4-1 to 4-n, respectively. And each of the interrupt right control circuits 3-1 to 3-n outputs an interrupt right transfer signal ID to a corresponding one of the interrupt right control circuits 3-2 to 3-1 of the next processor modules 1-2 to 1-1 a predetermined time after the interrupt right control circuit has the interrupt right when it does not receives an interrupt request from the interrupt controller 4-1 to 4-n, respectively. Each of the interrupt right control circuits 3-1 to 3-n which outputs the interrupt right transfer signal ID loses the interrupt right, and the corresponding one of the interrupt right control circuits 3-2 to 3-1 which receives the interrupt right transfer signal ID acquires the interrupt right. Note that the interrupt right control circuit 3-n of the last stage outputs the interrupt right transfer signal ID to the interrupt right control circuit 3-1 of the first stage.

Numerous aspects of this cannot be reconciled with the claimed invention.

Docket No.: 60751.300101

Patent

The invention of Yamada requires "chains" of respective sets of connected devices, with each chain ending in a respective processor (see e.g., FIG. 1). Thus, each CPU (2-1 to 2-n) of Yamada has a corresponding interrupt controller (4-1 to 4-n). In contrast, Applicant's independent claims 1 and 12 recite "*an interrupt controller*" (singular) and independent claims 39-40 recite "*a peripheral controller*" (also singular). Independent claims 23 and 31 are herein amended by moving into them the subject matter of their dependent claims 24 and 32 (with remarks provided below as to why this subject matter is not taught or reasonably suggested by Yamada).

Each of the CPU and interrupt controller pairings of Yamada further has a corresponding interrupt right control circuit (e.g., 2-1 with 3-1 and 4-1, ... 2-n with 3-n and 4-n) and the overall invention of Yamada cannot function without such. In contrast, the claimed invention does not require or use such an arrangement, our claims recite no such element or limitation.

Each interrupt right control circuit (3-1 to 3-n) outputs an interrupt right transfer signal (ID) to its subsequent interrupt right control circuit (e.g., 3-2 to 3-3 and 3-n to 3-1, with the latter particularly noted by Yamada). In contrast, the claimed invention does not have elements or steps, or equivalents, that operate in this passing-off or circular buffer-like manner. The claimed invention instead provides parallel interrupt rights control, such that even multiple processors can service the very same interrupt request.

Specifically, the Action states:

As for claims 1, 12, 23, 31 and 39-40, Yamada teaches a computer system having an interrupt handling apparatus ... comprising ... [1] an interrupt controller that is capable of handling the interrupts from the one or more hardware devices and capable of [2] independently generating a low priority interrupt signal and a high priority interrupt signal for each processor (see figure 1, interrupt controller 4-1 to 4-n and column 8 line 66 to column 8 line 64, [SIC, col. 7, ln. 66] wherein each I/O devices generates an interrupt to the each interrupt controller, [3] the interrupt controller then compares the priority interrupt level of each I/O device with the interrupt level set in the corresponding interrupt controller if the interrupt level is higher than the priority level set the interrupt controller routes that interrupt to the corresponding CPU and if the interrupt level is lower than the priority level set the interrupt controller negated that interrupt signal); (reference numbering added)

Point [1] has already been shown to be error above. Yamada teaches a plurality of interrupt controllers necessarily equaling its CPUs in quantity. Contrary to point [2], Yamada does not teach "*a low priority interrupt signal and a high priority interrupt signal for each processor.*"

Docket No.: 60751.300101

Patent

Anything resembling "prioritizing" in Yamada is performed by its interrupt mask circuits (5-1 to 5-n) and its interrupt right control circuits (3-1 to 3-n) chained together by its interrupt right transfer signals (ID). This is not a high/low prioritizing, and it is performed with circuits having no equivalents in the claimed invention. With respect to point [3], this is not what any of claims 1, 12, 23, 31, or 39-40 recites. Additionally, it appears that Applicant's providing both high and low interrupt signals concurrently has been confused with comparing high and low signals concurrently.

The Action next states:

the interrupt controller further comprising [4] an enable device that is capable of enabling the interrupt signal independently for each processor and a priority device that is capable of assigning a priority to each interrupt signal destined for any processor wherein a particular interrupt signal is capable of being routed to either processor and is capable of being assigned a low priority or high priority (see figure 1, column 7 line 66 to column 8 line 64, wherein the interrupt controller compares the request priority level from the I/O devices with the priority set in the interrupt controller, [5] if the request priority level is higher than the priority level set, the interrupt controller routes to corresponding CPU other wise if the request priority level is lower than the priority level set, the interrupt controller negated the request). (additional reference numbering added)

As point [4], the claimed invention has "an enable device that is capable of enabling the interrupt signal independently for each processor," yet Yamada teaches a plurality of interrupt mask circuits (5-1 to 5-n), interrupt controllers (4-1 to 4-n), and interrupt right control circuits (3-1 to 3-n) using sets of each necessarily equaling its CPUs in quantity. With respect to point [5], this is not what any of claims 1, 12, 23, 31, or 39-40 recite. Furthermore, it again appears that providing and comparing highs and lows have been confused. Although not explicitly recited in the claims at issue here, Applicant teaches a register based "mapping" of interrupts rather than a high verses low comparison as this point implies and as Yamada does teach.

The § 102(b) rejections as applied to claims 2 and 13:

We first urge that these should be allowable for at least the same reasons in the remarks above, for parent claims 1 and 12.

The Action here states:

... Yamada teaches first processor enable device that is capable of enabling the interrupt signals for the first processor and the a second processor enable device that is capable of enabling the interrupt signals for the second processor (see figure 1, I/O devices generates IRI to IRm requests and the

Docket No.: 60751.300101

Patent

interrupt controller examined these request in order to route each individual request signal to a corresponding CPU).

However, the claimed invention here has *"an enable device [singular] that is capable of enabling the interrupt signal independently for each processor"* and within that device further has respective processor enable devices (sub-deviccs) to enable the interrupt signals for the respective processors. What has been missed here is that the enable device enables interrupt signals independently, and the processor enable sub-deviccs then "drive" the FIQ and IRQ signals to their respective processor. How this differs from Yamada can be seen in Applicant's FIG. 4 (elements 192 and 194) and FIG. 5 (elements 246, 248 for FIQ.P and IRQ.P distinct from elements for FIQ.S and IRQ.S).

Yamada's general scheme is quite different, particularly by its use of its interrupt right control circuits (3-1 to 3-n) connected with ID signals. The ID signals are used to intercommunicate between its sets or "channels" of an interrupt mask circuit (5-x), interrupt controller (4-x), and interrupt right control circuit (3-x). This intercommunication occurs only at the very last, before its processors (2-x). This works just as Yamada states *"transferring the interrupt right"* (see e.g., Abstract), that is, passing it off until one channel can finally handle it.

In contrast, as can be appreciated from Applicant's FIG. 4-6, the use of an overall single enable device (recited in the claims), provides a number of advantages. For instance, it allows all interrupt requests to be visible to all the processors, although only the few that are assigned to any specific processor are serviced directly. Since each processor can read the status of the interrupt requests assigned to the others, it can "know" and "plan" for the tasks that maybe required for the other processors. In collaborative computing embodiments, this works as a sharing mechanism, allowing the work-load us to be distributed uniformly across processors. Directly, this ensures that work is completed sooner. Indirectly, the inventor has further found that this reduces power consumption.

The § 102(b) rejections as applied to claims 3-4, 14-15, 24-25 and 32-33 (and 23 and 31):

We first urge that these should be allowable for at least the same reasons in the remarks above, for parent claims 1, 12, 23, and 31.

We next note that dependent claims 24 and 32 are canceled as redundant, in view of their subject matter being moved into parent claims 23 and 31. And we submit that the subject matter

Docket No.: 60751.300101

Patent

of claims 24 and 32 makes now amended claims 23 and 31 allowable under the rationale of the rest of this section.

Turning now to the Action, it here states "... *Yamada teaches each enable device comprises one or more flip-flops connected together with one or more logical gates (see figure 4 and column 4 lines 55-67).*" Superficially this seems to read on our claims 3 and 14. However, virtually all digital devices today include flip-flops and logical gates. More determinative is that these common components are used in our overall enable device, and we have show above that Yamada does not teach or reasonably suggest this device.

With respect to claims 4, 15, 24-25, and 32-33, these all recite a register element or the use of such. The Action fails to state any argument whatsoever that Yamada teaches or suggests this.

The § 102(b) rejections as applied to claims 5, 16, 27, 30, 35 and 38:

We first urge that these should be allowable for at least the same reasons in the remarks above, for parent claims 1, 12, 23, and 31.

The Action here states:

... Yamada teaches the priority device further comprises a first processor priority device that is capable of assigning a priority to the interrupt signals for the first processor and a second processor priority device that is capable of assigning a priority to the interrupt signals for the second processor (see column 3 lines 32-38).

Respectfully, col. 3, ln. 32-38 provides no support for the broad assertion stated. Yamada here merely states that "each of the interrupt controllers (4-1 to 4-n) performs control [including] priority control." First, this reminds us that Yamada teaches the use of multiple interrupt channels, one per processor, and this is markedly different than our claimed invention.

Second, using parent claim 1 and dependent claim 5 as an example, the claimed invention here has "*a priority device [singular] that is capable of assigning a priority to each interrupt signal destined for any processor wherein a particular interrupt signal is capable of being routed to either processor.*" As can be seen in its FIG. 1, Yamada uses plural interrupt controllers (4-1 to 4-n). One of ordinary skill in the art can readily appreciate that Yamada is not capable of routing interrupt signals assigned by priority to respective processors.

In contrast, our use of an overall single priority device (recited in the subject claims), provides advantages such as allowing interrupt requests to be serviced directly by a specific

Docket No.: 60751.300101

Patent

processor, processor "knowing" and "planning" for inter-task requirements, improved collaborative computing, work-load distribution, etc.

The § 102(b) rejections as applied to claims 6 and 17:

We first urge that these should be allowable for at least the same reasons in the remarks above, for parent claims 1 and 12.

The Action here states "...*Yamada teaches each priority device comprises one or more flip flops connected together with one or more logical gates (see figure 4 and column 4 lines 55-67).*" Again, this superficially seems to read on our claims. However, this is not determinative of anything, since virtually all digital devices today include flip-flops and logical gates. Yamada does not teach or reasonably suggest using such in the manner of the claimed invention. Furthermore, FIG. 4 and col. 4, ln. 55-67 merely show and discuss the use of flip-flops to latch Yamada's IE and IA signals going into its processors. Any prioritization, using flip-flops or otherwise, must take place before these flip-flops.

The § 102(b) rejections as applied to claims 7-9, 18-20, 26, 28-29, 34 and 36-37:

We first urge that these should be allowable for at least the same reasons in the remarks above, for parent claims 1, 12, 23 and 31.

The Action here states:

... Yamada teaches each priority device further comprises a register containing an interrupt priority signal that is fed into the logical gates in order to determine the priority for the interrupt signals of the particular processor (see figure 1, I/O devices generates IRI to IRm requests, the interrupt controllers determines priority level by comparing its set priority level if the request priority level is higher then that request routes to corresponds CPU otherwise negated the request).

Respectfully, this asserts that Yamada teaches a register (apparently singular) -- yet then fails to provide any support for such an assertion. FIG. 1 of Yamada nowhere shows a register, and we fail to perceive any relevance to what its I/O devices generate or how its multiple interrupt controllers determine priority.

The § 102(b) rejections as applied to claims 10-11 and 21-22:

Docket No.: 60751.300101

Patent

We first urge that these should be allowable for at least the same reasons in the remarks above, for parent claims 1 and 12.

The Action here states "... *Yamada teaches the interrupt controller further comprises a device for multiplexing the one or more hardware interrupt signals with one or more software interrupt signals so that the processors are capable of being interrupted by hardware and software (see figure 1 and column 2 lines 51-64).*"

Respectfully, the cited portions of Yamada appear to provide no support whatsoever for this. If we have somehow failed to perceive the Examiner's rationale here we ask him to explain it so that Applicant is provided a reasonable opportunity to reply.

CONCLUSION

Applicant has endeavored to put this case into complete condition for allowance. It is thought that the §102 rejections are addressed by amendment or else shown to be unfounded on the prior art references cited. Applicant therefore asks that all objections and rejections now be withdrawn and that allowance of all claims presently in the case be granted.

Intellectual Property Law Offices
1901 S. Bascom Ave., Suite 660
Campbell, CA 95008

Telephone: 408.558.9950
Facsimile: 408.558.9960
E-mail: RRoberts@iplo.com

Respectfully Submitted,



Raymond E. Roberts
Reg. No.: 38,597

Customer No.:

32112